



IPW

Docket No.: 566.43530X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re the Application of:

Noboru MORISHITA et al.

Serial No. 10/783,018

Filed: February 23, 2004

For: REMOTE COPY SYSTEM

SUPPLEMENTAL PETITION TO MAKE SPECIAL

June 17, 2005

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Supplemental to the Request for Reconsideration filed on May 18, 2005,
in view of the meeting between Mr. Brundidge and Mr. Laufer held on June 9,
2005 clarifying issues related to the granting of Petitions to Make Special,
Applicants submit the following additional remarks.

It is submitted that the cited references, whether considered alone or in combination, fail to disclose or suggest the invention as claimed. In particular, the cited references, at a minimum, fail to disclose or suggest in combination with the other limitations recited in the claims:

a first feature of the present invention as recited in independent claim 1 wherein when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, said second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, said second storage system stores write data received from the host computer and the data received from said first storage system and stored in said first primary volume, and sends data determined based on said management information out of the data stored in said second primary volume, to said third storage system;

a second feature of the present invention as recited in independent claim 8 wherein when the data stored in said first primary volume is migrated to said second primary volume, said host computer sends a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, said second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, said second storage system stores write data received from the host computer and the data

received from said first storage system and stored in said first primary volume, into said second primary volume, and sends data determined based on said management information out of the data stored in said second primary volume, to said third storage system;

a third feature of the present invention as recited in independent claim 9 wherein when the data stored in said first primary volume is migrated to said second primary volume, said management apparatus sends an access request received from the host computer and destined to a primary volume, to said second storage system, said second storage subsystem receives, from said first storage subsystem, management information for identifying data to send to said second storage system, said second storage system stores write data received from the host computer and the data received from said first storage system and stored in said first primary volume, into said second primary volume, and sends data determined based on said management information out of the data stored in said second primary volume, to said second storage system; and

a fourth feature of the present invention as recited in independent claim 10 recites a code that said third storage system stores said write data received from said first storage system into said secondary volume, when data stored in said first primary volume is migrated to a second volume owned by a second storage system, a code that said first storage system sends management information for identifying data to send said third storage system, to said second storage system, a code that said second storage system stores write data received from the host computer and data received from said first storage system, into said

second primary volume, and a code that said second storage system sends data identified based on said management information, out of data stored in said second primary volume, to said third storage system.

To the extent applicable to the present Petition, Applicants submit that although the distinguishing feature(s) may represent a substantial portion of the claimed invention, the claimed invention including said feature(s) and their inter-operation provides a novel storage system and system and method related to or implemented in or by said storage system not taught or suggested by any of the references of record.

The references considered most closely related to the claimed invention are briefly discussed below:

U.S. Patent No. 5,673,382 (Cannon et al.) discloses a data processing system providing disaster recovery and using a client-server configuration provides a method and apparatus for reclaiming off-site storage volumes without requiring the volumes to be mounted or transported to the primary site. The system includes primary volumes and resident copy volumes at the primary site and off-site storage volumes for disaster recovery at the remote site. The system first determines which off-site storage volumes are eligible for reclamation. For each eligible off-site volume, the system locates, within the primary site, a primary copy of every file remaining on the off-site volume that is still required for disaster recovery. These primary copies are then copied to a resident recovery volume at the primary storage site. The eligible off-site volume is then marked empty, and the resident recovery volume is classified as an off-site volume. The

resident recovery volume is then transported to the remote storage site and the empty off-site volume is returned to the primary storage site to be reused as a resident copy volume. However, Cannon et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Cannon et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,253,295 (Beal et al.) discloses a data storage system including first, second, third and fourth virtual storage volumes, the first and second virtual storage volumes controlled to enable a mirroring and synchronization of data therebetween. The data storage system further includes a first processor that is responsive to a "copy pair" command or its equivalent to copy a directory from the first virtual storage volume to a directory of the third virtual storage volume. The first processor also dispatches a copy command, or its equivalent, to copy the directory of the second virtual storage volume to the directory of the fourth storage volume. A second processor is responsive to the copy command to cause a copying of the directory of the second virtual storage volume to the directory of the fourth virtual storage volume. Thus, the third and fourth virtual storage volumes are controlled to include directories which duplicate directories of the first and second virtual storage volumes. However, Beal et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of

said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Beal et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,754,792 (Nakamura et al.) discloses a storage system data duplication method for copying data of a plurality of logical volumes possessed by a first storage system to a second storage system includes the steps of: (a) copying the data of the plurality of logical volumes to the second storage system; (b) interrupting copying of the data of the plurality of logical volumes to the second storage system; (c) starting copying of data of one or more of the plurality of logical volumes to the second storage system; and starting copying of data of ones of the plurality of logical volumes other than the one or more logical volumes to the second storage system as delayed from the step (c), shortening a time taken until completion of data duplication on recreation of a group of paired logical volumes in a volume group after remote copy is temporarily stopped. However, Nakamura et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access

request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Nakamura et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2003/0033494 A1 (Fujibayashi et al.)

discloses, in one embodiment, a new storage system is connected to an old storage system destined for replacement. Prior to starting online data migration, remote copy configuration information from the old storage system is migrated to the new storage system. The aforementioned allows either the new storage system or the old storage system, depending on configuration, to maintain an

operational remote copy function with another storage system during the online data migration. The old and new storage systems can be primary storage systems located in a local site. The old and new storage systems can also be secondary storage systems located in a site remotely located from a local site. However, Fujibayashi et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Fujibayashi et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent

claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2003/0126388 A1 (Yamagami) discloses techniques for managing storage based replication. Specific embodiments provide the capability to make multiple remote copies of information without the necessity of copying information for each pair, for example. One or more remote mirror pairs can be created for disaster recovery, testing, or other purposes. However, Yamagami et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Yamagami et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described

third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2003/0177321 (Watanabe) discloses a method and apparatus are provided for enhancing the performance of storage systems to allow recovery after all types of suspensions in remote copy operations. Data is synchronized after an interruption in transfer between a first storage volume of a primary storage system and a first storage volume of a secondary storage system which also includes a second storage volume. After the interruption is detected, at the primary storage system, a record is provided of the data written onto the first storage volume of the primary storage system, and at the secondary storage volume a record is provided of the data written onto the first storage volume of the secondary storage system. Then, at least a partial copy of the record of the data written onto the first storage volume of the primary storage system is written onto the second storage volume. Using the copy, the first storage volume of the secondary storage system is synchronized with the second storage volume of the secondary storage system. However, Watanabe et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management

information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Watanabe does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0078399 A1 (Tabuchi et al.) discloses in order to realize asynchronous type system assuring the consistency of data with the function of disk subsystems without the need of introducing new software to a host unit and without the deterioration of the performance of a main center, in a remote copy system which copies the data to the disk subsystems of the remote center for duplicating the data in the disk subsystems of the main center, the disk subsystems give serial numbers and times to the data together with writing said data to the storage devices in the disk subsystem and transfer said data to the other disk subsystems, and the other disk subsystems arrange

the two or more data in the sequence of the serial numbers, decide the oldest time among the latest time given to each of the disk subsystems communicating among the disk subsystems and the data given with the time not later than the decided oldest time are the objects of writing to each of the disk storage devices. However, Tabuchi et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Tabuchi et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent

claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0107315 A1 (Watanabe et al.)

discloses in a storage system, a logical volume is divided into a plurality of small areas called logical storage devices and only such an area designated by a user is made an object of remote copying or migratory copying. Also, there is provided a unit for forming a logical volume from any logical storage device of any RAID group. Thereby, the reduction of the deterioration in performance at the time of remote copying, the reduction of a storage area to be assigned to a secondary site as a copy destination, the shortening of a time required for migratory copying and the assignment of any area of a logical volume on any RAID group are enabled. However, Watanabe et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host

computer and destined to a primary volume, to said second storage system. More particularly, Watanabe et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0128363 A1 (Yamagami et al.)

discloses a method for handling a remote copy request in a distributed storage includes providing a plurality of primary volumes within a primary storage system that is coupled to a primary host via a first network, the primary storage system being coupled to a secondary storage system via a second network. A first request is selected from a plurality of requests placed in a queue based on priority information associated with the requests. A first path group is selected from one or more path groups that could be used to transmit the request. The first request is transmitted to the secondary storage system using the first path group, the secondary storage system including a plurality of secondary volumes that are paired to the plurality of primary volumes. However, Yamagami et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage

system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Yamagami et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0148443 A1 (Achiwa) discloses a storage system is composed of storage apparatuses each separately installed at three sites for the purpose of disaster recovery, wherein data stored in one of the storage apparatuses designated as a replication source is also managed in one of the other two storage apparatuses designated as a replication destination, such that in the event of a failover execution request from an information processing apparatus, data stored in the storage apparatus designated as the replication source is managed also in the storage apparatus of the other two

storage apparatuses designated as a spare. However, Achiwa, at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Achiwa et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0153719 A1 (Achiwa et al.) discloses an information processing system includes three storage apparatuses installed on three sites for the purpose of disaster recovery. The three storage

apparatuses can be accessed by three information processing apparatuses, respectively, that are configured in a cluster. One of the storage apparatuses is set as a replication source and the other two storage apparatuses are set as replication destinations. The two storage apparatuses that are set as replication destinations manage a copy of data stored in the storage apparatus that is set as the replication source. In association with an execution of a failover on the information processing apparatus side, settings of the storage apparatuses as the replication source and replication destinations are automatically changed such that the storage apparatus that is set as the replication source would function as a replication destination, and one of the storage apparatuses that are set as the replication destinations would function as a replication source. However, Achiwa et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a

primary volume, to said second storage system. More particularly, Achiwa et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0193660 A1 (Gagne et al.) discloses a data storage facility for transferring data from a data altering apparatus, such as a production data processing site to a remote data receiving site. The data storage facility includes a first data store for recording each change in the data generated by the data altering apparatus. A register set records each change on a track-by-track basis. A second data store has first and second operating modes. During a first operating mode the second data store becomes a mirror of the first data store. During a second operating mode the second data store ceases to act as a mirror and becomes a source for a transfer of data to the data receiving site. Only information that has been altered, i.e., specific tracks that have been altered, are transferred during successive operations in the second operating mode. Commands from the local production site initiate the transfers between the first and second operating modes. However, Gagne et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus

transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Gagne et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0205294 A1 (Nakayama et al.)

discloses in a storage device controlling apparatus which includes: a plurality of channel controllers having a circuit board on which are formed a file access processing section receiving from an information processing apparatus requests to input and output data in files as units via a network and an I/O processor outputting to a storage device I/O requests corresponding to the requests to input

and output data; and a disk controller executing input and output of data into and from the storage device in response to the I/O requests sent from the I/O processors, at least one of the channel controllers receives data specifying an assignment of a logical volume to the channel controller, the data being sent from the information processing apparatus, and stores the received assignment. However, Nakayama et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Nakayama et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent

claim 10, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0230859 A1 (Cochran et al.)

discloses a disaster recovery system with sequenced cascaded resynchronization comprises a plurality of data centers and a distributed control system. The individual data centers comprise a communication interface, a data storage, and a controller. The distributed control system is distributed and executable in the controllers of the plurality of data centers, and is capable of coordinating operations via the communication interfaces of the plurality of data centers to resynchronize a plurality of communication links between data center pairs of the plurality of data centers. The communication links including at least one synchronous link and at least one asynchronous link. However, Cochran et al., at a minimum, fails to disclose or suggest when the data stored in said first primary volume is migrated to said second primary volume, said network apparatus transfers an access request issued from the host computer and destined to said first primary volume, to said second primary volume, or the second storage system receives, from said first storage system, management information for identifying data to send to said third storage system, or sending data determined based on said management information out of the data stored in said second primary volume, to said third storage system, or a host computer sending a write request issued from an application program executed by the host computer and destined to a primary volume, to said second primary volume of said second storage system, or a management apparatus sending an access


request received from the host computer and destined to a primary volume, to said second storage system. More particularly, Cochran et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims.

Therefore, since the cited references fail to disclose or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 8, the above described third feature of the present invention as recited in independent claim 9, and the above described fourth feature of the present invention as recited in independent claim 10, in combination with the other limitations recited in each of the independent claims, it is submitted that all of the claims are patentable over the cited references whether said references are taken individually or in combination with each other.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of Mattingly, Stanger, Malur & Brundidge, P.C., Deposit Account No. 50-1417 (referencing attorney docket no. 566.43530X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Frederick D. Bailey
Registration No. 42,282

FDB/sdb
(703) 684-1120